The OSG Resource Selection Service (ReSS)

Gabriele Garzoglio
Fermilab, Computing Division
March 13, 2007

The Resource Selection Project

- The Resource Selector Service implements cluster-level Workload Management on OSG.
- The project started in Sep 2005
- Sponsors
  - DZero contribution to the Common Project
  - FNAL-CD (30% FTE Gabriele, 50% FTE Tanya)
- Collaboration of the Sponsors with
  - OSG (TG-MIG, ITB, VDT / John Weigand)
  - CEMon gLite Project (INFN)
  - FermiGrid
  - Glue Schema Group

The Resource Selection Service Motivations / Deliverables

- A Resource Selector allows…
  - …expressing requirements on the resources in the job description
  - …the user to refer to *abstract* characteristics of the resources in the job description
- The Resource Selection Project has two major goals
  1. Enable OSG resource usage by DZero. Jobs are prepared and data is handled by the SAM-Grid.
  2. Develop and deploy a Resource Selection Service that VOs with requirements on job management similar to DZero can use.

Resource Selection Example

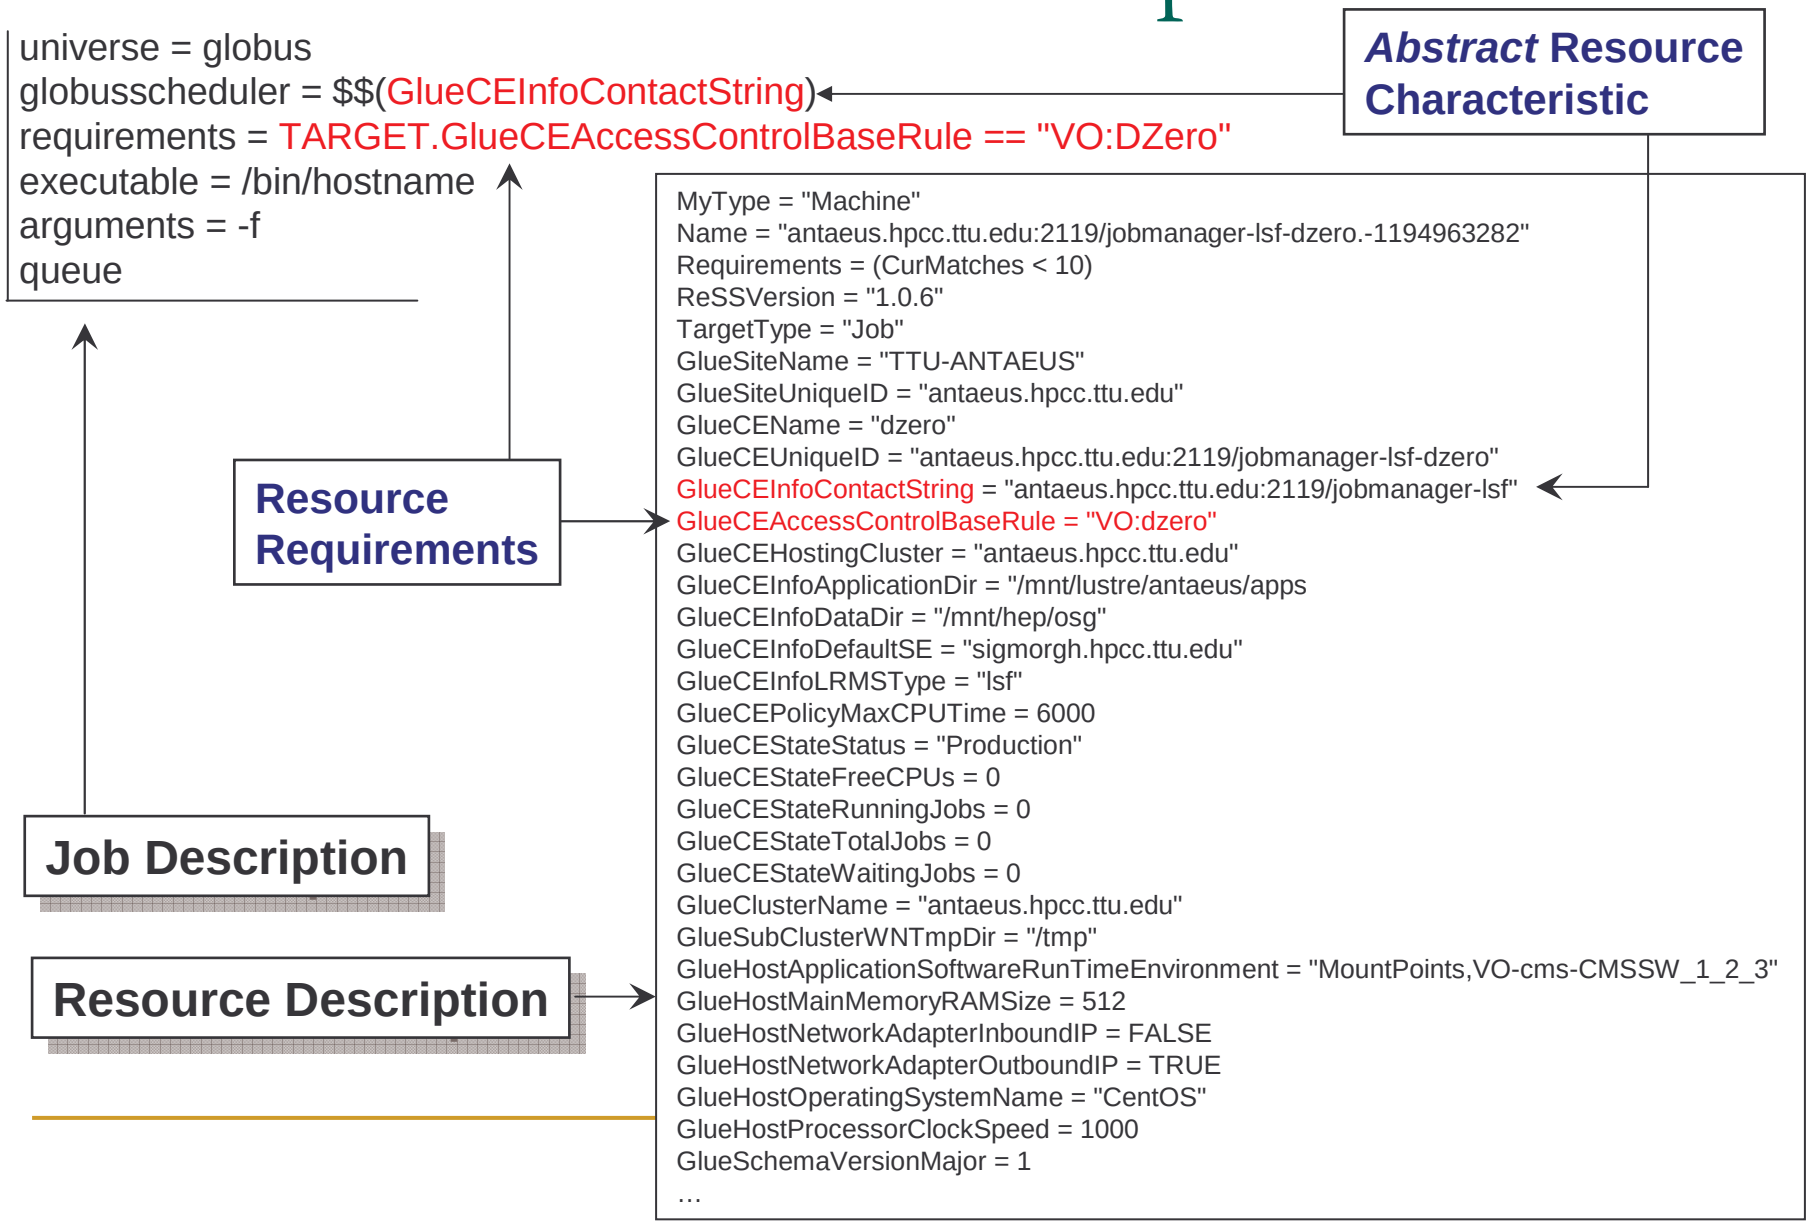

The Resource Selection Service Architecture

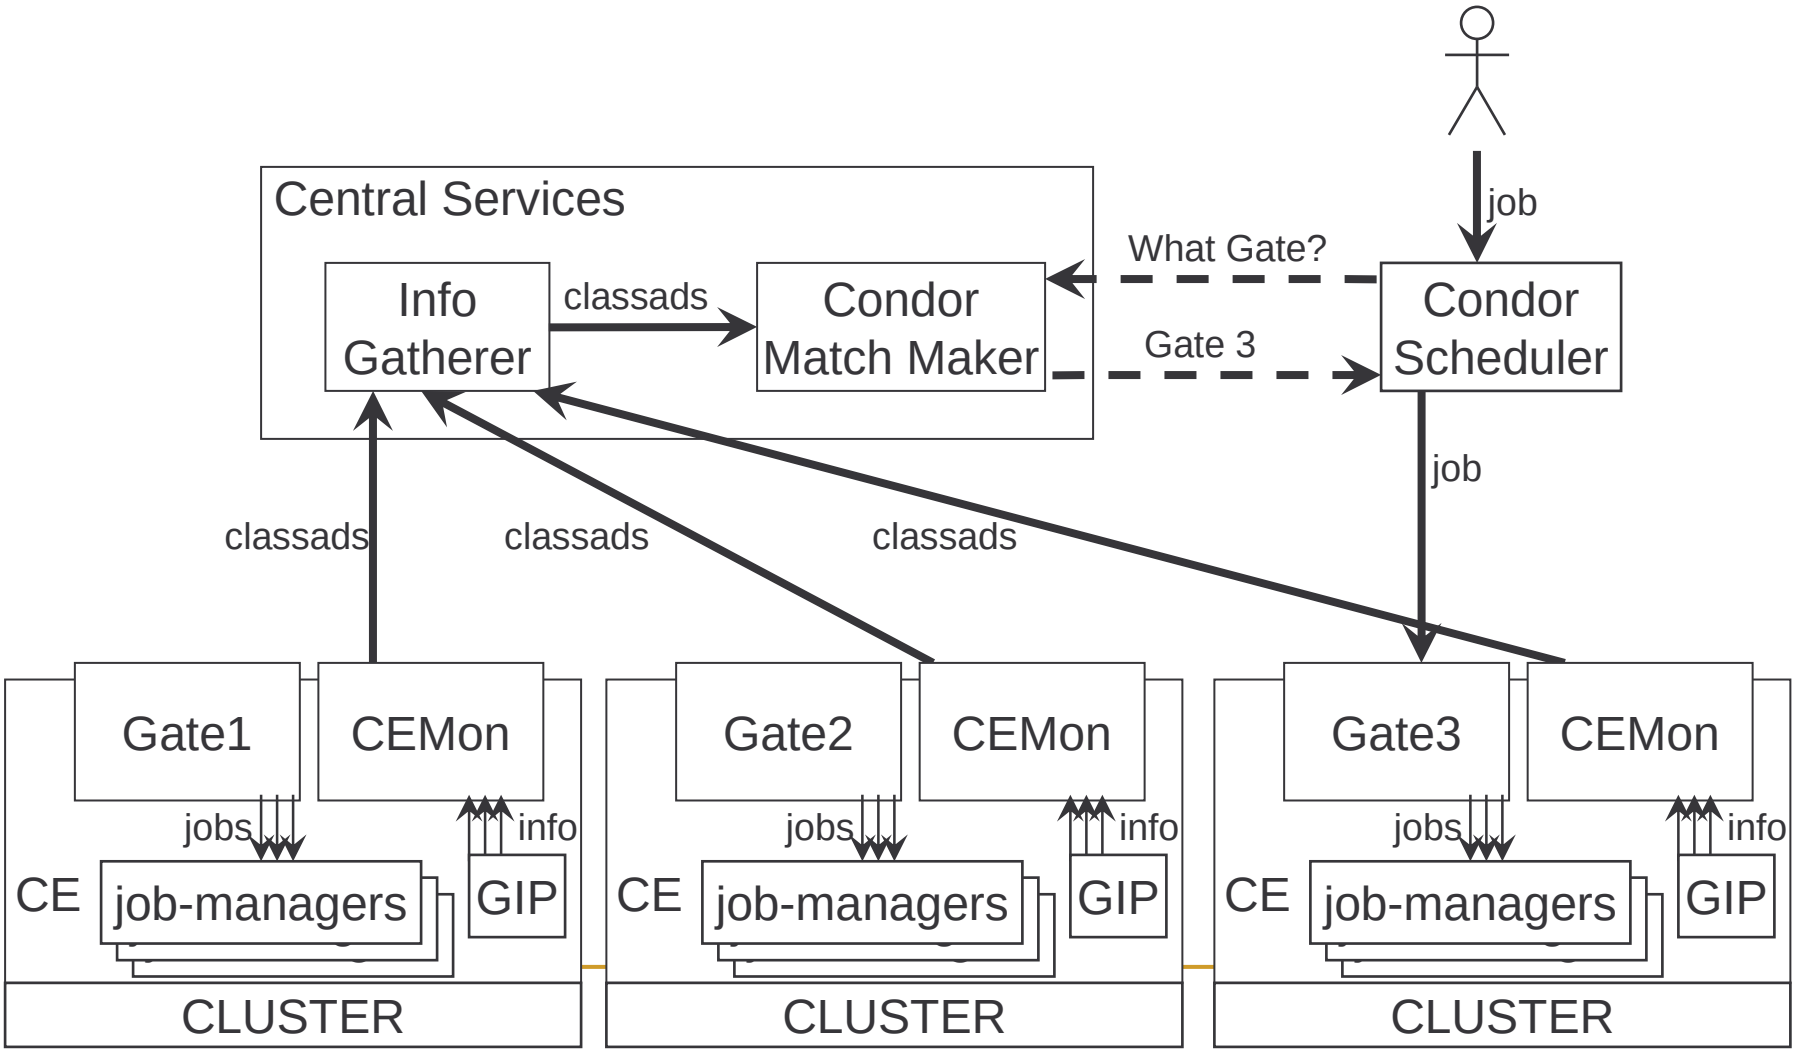

ReSS Validation

- **Validated that requirements of DZero are met by the ReSS central services**
  - https://twiki.grid.iu.edu/twiki/bin/view/ResourceSelection/ReSSValidationTest

Jobs Submitted = 1 job/sec for hour.
Total Jobs Submitted = 3600
First Job Matched = 9/8/2006 16:33:00
Last Job Matched = 9/9/2006 02:05:53
Resources Satisfying Jobs = 2 (1800 jobs per resource)
Total Number Of Resources = 426
Max Jobs Matched Per Negotiation Cycle Per Resource = 10
Total Jobs Matched In One Negotiation Cycle = 20
Longest Negotiation Cycle: 2 sec
Shortest Negotiation Cycle: 0 sec
Average Negotiation Cycle: 0.772222222222 sec

- **Investigated the impact on resources (load, mem, …) of CEMon at OSG CEs**
  - https://twiki.grid.iu.edu/twiki/bin/view/ResourceSelection/CEMonPerformanceEvaluation

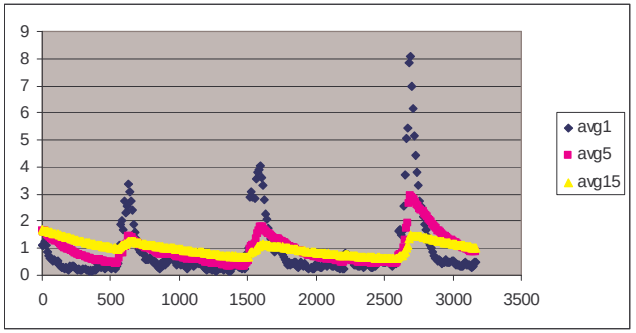

- **US CMS studied the scalability of ReSS central services for US CMS requirements**
  - https://twiki.grid.iu.edu/twiki/bin/view/ResourceSelection/ReSSEvaluationByUSCMS

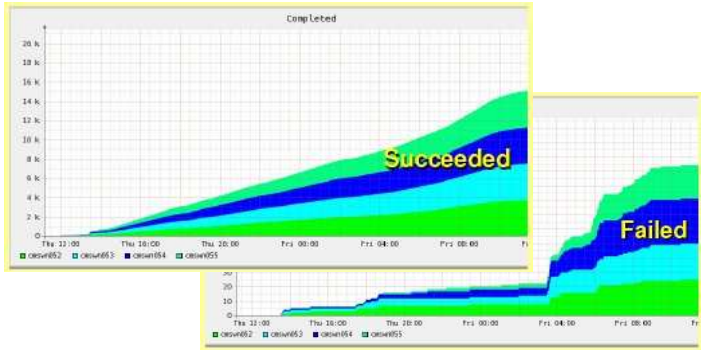

Status

- Development is mostly done
  - We may still add SE to the resource selection process
- Integration of ReSS with Fermigrid is done
- Assisting Deployment of ReSS on Production OSG
  - Worked with ITB since May 06, targeting deployment for Summer 06
  - Validation process very slow: OSG 0.6.0 released on Mar 07.
- Using ReSS on SAM-Grid / OSG for DZero data reprocessing for the available sites
  - However, the delay in OSG deployment makes operations difficult (keeping “right” amount of idle jobs at sites)
- Working with OSG VOs to facilitate ReSS usage

Current Deployment

Site	Gatekeeper	2007-03-05-14-00-08	2007-03-05-18-00-09	2007-03-06-00-00-13	2007-03-06-06-00-10	2007-03-06-12-00-09	2007-03-06-18-00-07	2007-03-07-00-00-14	2007-03-07-06-00-08	2007-03-07-12-00-11	2007-03-07-18-00-09
CornellLEPP	lnx6211.lns.cornell.edu:2119/jobmanager-sge	Down	Down	Down	Down	Down	Down	Down	Down	Down	Up
NERSC-VM-VTB0	osp-vtb00.nersc.gov:2119/jobmanager-sge	Up	Down	Up	Up	Down	Down	Down	Down	Down	Down
TTU-ANTAEUS	antaeus.hpcc.ttu.edu:2119/jobmanager-lsf	Down	Down	Down	Down	Down	Down	Down	Down	Down	Down
UCRHEP	top.ucr.edu:2119/jobmanager-condor	Down	Down	Down	Down	Down	Down	Down	Down	Down	Up
UCTier3	uct3-edge6.uchicago.edu:2119/jobmanager-pbs	Up	Up	Up	Up	Up	Up	Up	Up	Down	Up
UVaHEP-T3	osg-hep.phys.virginia.edu:2119/jobmanager-pbs	Down	Down	Down	Down	Down	Down	Down	Down	Down	Up
Vanderbilt	vmpg01.vampire:2119/jobmanager-pbs	Down	Down	Down	Down	Down	Down	Down	Down	Down	Up
cmsosgce.fnal.gov:2119/jobmanager-condor		Up	Up	Up	Up	Up	Up	Up	Up	Up	Up
red.unl.edu:2119/jobmanager-pbs		Up	Up	Up	Up	Up	Up	Up	Up	Up	Up
sammy.fnal.gov:2119/jobmanager-condor		Up	Up	Up	Up	Up	Up	Up	Up	Up	Up
stitch.oscer.ou.edu:2119/jobmanager-condor		Up	Up	Up	Up	Up	Up	Up	Up	Up	Up

Click here for live URL

Remaining Tasks for the Project

- Assist with OSG deployment (i.e. CEMon at sites)
- Assist OSG VOs (e.g. Engagement) to use ReSS
- Integrate ReSS with GlideIn Factory
- Check with collaborators if they are interested in SE support
  - …one of the last development activities on the table today
- Assist OSG with Truth-In-Advertisement (GIP)
- Move project from devel. to maintenance
  - estimated effort reduction: from 0.8 FTE to 0.25 FTE
- Maintain CEMon in VDT reasonably up to date

Conclusions

- ReSS Project is naturally moving from development to maintenance
- We are still involved in integration and supporting activities
- More info at http://osg.ivdgl.org/twiki/bin/view/ResourceSelection/